

CLAIMS

1. A method for managing a data storage system that includes first and second storage subsystems, the method comprising:

copying data stored on the first storage subsystem to the
5 second storage subsystem in an asynchronous mirroring process;

maintaining a record on the second storage subsystem,
indicative of locations at which the data have been updated on
the first storage subsystem and have not yet been copied to
the second storage subsystem;

10 receiving at the second storage subsystem, from a host
processor, a request to access the data stored at a specified
location on the data storage system;

if the specified location is included in the record,
sending a message from the second storage subsystem to the
15 first storage subsystem requesting a synchronous update of the
data at the specified location; and

providing the data stored at the specified location from
the second storage subsystem to the host processor after
receiving at the second storage subsystem a response to the
20 message from the first storage subsystem.

2. The method according to claim 1, and comprising providing
the data stored at the specified location from the second
storage subsystem to the host processor in response to the
25 request without sending the message if the specified location
is not included in the record.

3. The method according to claim 1, and comprising removing the specified location from the record on the second storage subsystem upon receiving the response from the first storage subsystem.

5

4. The method according to claim 1, wherein the locations included in the record maintained on the second storage subsystem are a superset of the locations at which the data have been updated on the first storage subsystem and have not
10 yet been copied to the second storage subsystem.

5. The method according to claim 4, wherein maintaining the record comprises, upon receiving an update to the data at a given location on the first storage subsystem, sending an
15 instruction from the first storage subsystem to the second storage subsystem, which causes the second storage subsystem to add a plurality of locations, including the given location, to the record.

20 6. The method according to claim 5, wherein sending the instruction comprises selecting the plurality of the locations to add to the record responsively to a prediction of the locations at which the data are expected to be updated subsequent to updating the data at the given location.

25

7. The method according to claim 5, wherein maintaining the record comprises maintaining a copy of the record on the first storage subsystem, and wherein sending the instruction

comprises determining whether to send the instruction responsively to the copy of the record.

8. The method according to claim 7, wherein determining
5 whether to send the instruction comprises deciding to send the instruction only if the given location is not included in the record.

9. The method according to claim 7, wherein maintaining the
10 record further comprises removing the specified location from both the record and the copy of the record responsively to receiving the response from the first storage subsystem at the second storage subsystem.

15 10. The method according to claim 4, wherein the response from the first storage subsystem comprises the data stored at the specified location if the specified location is one of the locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second
20 storage subsystem, and wherein the response otherwise comprises a status notification indicating that the data at the specified location on the second storage subsystem are synchronized with the data on the first storage subsystem.

25 11. The method according to claim 1, and comprising incorporating the data stored at the specified location in the response from the first storage subsystem to the second storage subsystem, and conveying the data from one or more

additional locations on the first storage subsystem to the second storage system responsively to the message from the second storage subsystem.

5 12. The method according to claim 11, wherein conveying the data from the additional locations comprises selecting the additional locations predictively based on the request from the host processor to access the data stored at the specified location.

10

13. The method according to claim 1, and comprising updating the data that are stored on the second storage subsystem responsively to an input from the host processor, and copying the updated data from the second storage subsystem to the
15 first storage subsystem in the asynchronous mirroring process.

14. The method according to claim 13, and comprising:

maintaining a further record on the first storage subsystem, indicative of the locations at the which the data
20 have been updated on the second storage subsystem and have not yet been copied to the first storage subsystem;

receiving at the first storage subsystem, from a further host processor, a further request to access the data stored at a further location on the data storage system;

25 if the further location is included in the record, sending a further message from the first storage subsystem to the second storage requesting a synchronous update of the data at the further location; and

providing the data stored at the further location from the first storage subsystem to the further host processor after receiving at the first storage subsystem the response from the second storage subsystem to the further message.

5

15. The method according to claim 14, wherein updating the data comprises receiving the data written by the host processor to be stored at a given location on the second storage system, and notifying the first storage system in a synchronous message that the data at the given location have
10 been updated if the given location is not included in the further record on the first storage subsystem.

16. The method according to claim 1, wherein copying the data
15 comprises transmitting the data between mutually-remote sites over a communication link between the sites.

17. The method according to claim 16, wherein the sites
comprise first and second mutually-remote sites, at which the
20 first and second storage subsystems are respectively located, and wherein the host processor from which the request is received at the second storage subsystem is located in proximity to the second site, and the method further comprises providing the data from the first storage subsystem to a
25 further host processor that is located in proximity to the first site.

18. The method according to claim 1, and comprising, upon occurrence of a failure in the second storage subsystem, configuring the first storage subsystem to provide the data directly to the host processor.

5

19. The method according to claim 1, wherein maintaining the record comprises marking respective bits in a bitmap corresponding to the locations that are included in the record.

10

20. The method according to claim 1, wherein the second storage subsystem comprises a computer workstation, and wherein maintaining the record and providing the data comprise maintaining the record and providing the data using an installable file system running on the computer workstation.

15

21. A data storage system, comprising first and second storage subsystems, which are arranged to store data,

wherein the first storage subsystem is arranged to copy

20

the data stored on the first storage subsystem to the second storage subsystem in an asynchronous mirroring process, and

wherein the second storage subsystem is arranged to maintain a record indicative of locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second storage subsystem, and

25

wherein the second storage subsystem is further arranged to receive a request from a host processor to access the data stored at a specified location on the data storage system, and

if the specified location is included in the record, to send a message to the first storage subsystem requesting a synchronous update of the data at the specified location, and to provide the data stored at the specified location to the host processor after receiving a response to the message from the first storage subsystem.

22. The system according to claim 21, wherein the second storage subsystem is arranged to provide the data stored at the specified location to the host processor in response to the request without sending the message if the specified location is not included in the record.

23. The system according to claim 21, wherein the second storage subsystem is arranged to remove the specified location from the record upon receiving the response from the first storage subsystem.

24. The system according to claim 21, wherein the locations included in the record maintained on the second storage subsystem are a superset of the locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second storage subsystem.

25. The system according to claim 24, wherein the first storage subsystem is arranged, upon receiving an update to the data at a given location, to send an instruction to the second storage subsystem, which causes the second storage subsystem

to add a plurality of locations, including the given location,
to the record.

26. The system according to claim 25, wherein the plurality
5 of the locations to add to the record are selected
responsively to a prediction of the locations at which the
data are expected to be updated subsequent to updating the
data at the given location.

10 27. The system according to claim 25; wherein the first
storage subsystem is arranged to maintain a copy of the
record, and to determine whether to send the instruction
responsively to the copy of the record.

15 28. The system according to claim 27, wherein the first
storage subsystem is arranged to send the instruction only if
the given location is not included in the record.

20 29. The system according to claim 27, wherein the first and
second storage subsystems are arranged to remove the specified
location from both the record and the copy of the record
responsively to receiving the response from the first storage
subsystem at the second storage subsystem.

25 30. The system according to claim 24, wherein the response
from the first storage subsystem comprises the data stored at
the specified location if the specified location is one of the

locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second storage subsystem, and wherein the response otherwise comprises a status notification indicating that the data at the specified location on the second storage subsystem are synchronized with the data on the first storage subsystem.

31. The system according to claim 21, wherein the first storage subsystem is arranged to send the response incorporating the data stored at the specified location, and to convey the data from one or more additional locations on the first storage subsystem to the second storage system responsively to the message from the second storage subsystem.

32. The system according to claim 31, wherein the first storage subsystem is arranged to select the additional locations predictively based on the request from the host processor to access the data stored at the specified location.

33. The system according to claim 21, wherein the second storage subsystem is arranged to update the data that are stored on the second storage subsystem responsively to an input from the host processor, and to copy the updated data to the first storage subsystem in the asynchronous mirroring process.

34. The system according to claim 33, wherein the first storage subsystem is arranged to maintain a further record

that is indicative of the locations at the which the data have been updated on the second storage subsystem and have not yet been copied to the first storage subsystem, and

5 wherein the first storage subsystem is arranged to receive a request from a further host processor to access the data stored at a further location on the data storage system, and if the further location is included in the further record, to send a further message to the second storage requesting a synchronous update of the data at the further location, and to
10 provide the data stored at the further location to the further host processor after receiving the response to the further message from the second storage subsystem.

35. The system according to claim 34, wherein the input from
15 the host processor comprises the data to be written to the given location on the second storage system, and wherein the second storage subsystem is arranged to notify the first storage system in a synchronous message that the data at the given location have been updated if the given location is not
20 included in the further record on the first storage subsystem.

36. The system according to claim 21, wherein the first and second storage subsystems are located at mutually-remote sites and are arranged to communicate over a communication link
25 between the sites.

37. The system according to claim 36, wherein the sites comprise first and second mutually-remote sites, at which the

first and second storage subsystems are respectively located,
and wherein the host processor from which the request is
received at the second storage subsystem is located in
proximity to the second site, and the first storage subsystem
5 is arranged to provide the data to a further host processor
that is located in proximity to the first site.

38. The system according to claim 21, wherein the first
storage subsystem is arranged, upon occurrence of a failure in
10 the second storage subsystem, to provide the data directly to
the host processor.

39. The system according to claim 21, wherein the record
comprises a bitmap, and wherein the second storage subsystem
15 is arranged to mark respective bits in the bitmap
corresponding to the locations that are included in the
record.

40. The system according to claim 21, wherein the second
20 storage subsystem comprises a computer workstation having an
installable file system running thereon, and wherein the
installable file system causes the second storage subsystem to
maintain the record and to provide the data to the host
processor.

25

41. A computer software product for use in a data storage
system including primary and secondary storage subsystems,
which are arranged to store data and include respective first

and second control units, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the first and second control units, cause the first control unit to copy the data stored on
5 the first storage subsystem to the second storage subsystem in an asynchronous mirroring process, and

wherein the instructions cause the second control unit to maintain a record indicative of locations at which the data have been updated on the first storage subsystem and have not
10 yet been copied to the second storage subsystem, and further cause the second control unit, upon receiving a request from a host processor to access the data stored at a specified location on the data storage system, if the specified location is included in the record, to send a message to the first
15 storage subsystem requesting a synchronous update of the data at the specified location, and to provide the data stored at the specified location to the host processor after receiving a response to the message from the first storage subsystem.

20 42. The product according to claim 41, wherein the instructions cause the second control unit to provide the data stored at the specified location to the host processor in response to the request without sending the message if the specified location is not included in the record.

25 43. The product according to claim 41, wherein the instructions cause the second control unit to remove the specified location from the record upon receiving the response from the first storage subsystem.

44. The product according to claim 41, wherein the locations included in the record maintained by the second control unit are a superset of the locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second storage subsystem.

45. The product according to claim 44, wherein the instructions cause the first control unit, upon receiving an update to the data at a given location, to send an instruction to the second storage subsystem, which causes the second control unit to add a plurality of locations, including the given location, to the record.

46. The product according to claim 45, wherein the plurality of the locations to add to the record are selected responsively to a prediction of the locations at which the data are expected to be updated subsequent to updating the data at the given location.

47. The product according to claim 45, wherein the instructions cause the first control unit to maintain a copy of the record, and to determine whether to send the instruction responsively to the copy of the record.

48. The product according to claim 47, wherein the instructions cause the first control unit to send the

instruction only if the given location is not included in the record.

49. The product according to claim 47, wherein the instructions cause the first and second control units to remove the specified location from both the record and the copy of the record responsively to receiving the response from the first storage subsystem at the second storage subsystem.

50. The product according to claim 44, wherein the response from the first storage subsystem comprises the data stored at the specified location if the specified location is one of the locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second storage subsystem, and wherein the response otherwise comprises a status notification indicating that the data at the specified location on the second storage subsystem are synchronized with the data on the first storage subsystem.

51. The product according to claim 41, wherein the instructions cause the first control unit to send the response incorporating the data stored at the specified location, and to convey the data from one or more additional locations on the first storage subsystem to the second storage system responsively to the message from the second storage subsystem.

52. The product according to claim 51, wherein the instructions cause the first control unit to select the

additional locations predictively based on the request from the host processor to access the data stored at the specified location.

5 53. The product according to claim 41, wherein the instructions cause the second control unit to update the data that are stored on the second storage subsystem responsively to an input from the host processor, and to copy the updated data to the first storage subsystem in the asynchronous
10 mirroring process.

54. The product according to claim 53, wherein the instructions cause the first control unit to maintain a further record that is indicative of the locations at the
15 which the data have been updated on the second storage subsystem and have not yet been copied to the first storage subsystem, and

wherein the instructions further cause the first control unit to receive a request from a further host processor to
20 access the data stored at a further location on the data storage system, and if the further location is included in the further record, to send a further message to the second storage requesting a synchronous update of the data at the further location, and to provide the data stored at the
25 further location to the further host processor after receiving the response to the further message from the second storage subsystem.

55. The product according to claim 54, wherein the input from the host processor comprises the data to be written to the given location on the second storage system, and wherein the instructions cause the second control unit to notify the first
5 storage system in a synchronous message that the data at the given location have been updated if the given location is not included in the further record on the first storage subsystem.

56. The product according to claim 41, wherein the first and
10 second storage subsystems are located at mutually-remote sites and are arranged to communicate over a communication link between the sites.

57. The product according to claim 56, wherein the sites
15 comprise first and second mutually-remote sites, at which the first and second storage subsystems are respectively located, and wherein the host processor from which the request is received at the second storage subsystem is located in proximity to the second site, and the instructions cause the
20 first control unit to provide the data to a further host processor that is located in proximity to the first site.

58. The product according to claim 41, wherein the instructions cause the first control unit, upon occurrence of
25 a failure in the second storage subsystem, to provide the data directly to the host processor.

59. The product according to claim 41, wherein the record comprises a bitmap, and wherein the instructions cause the second control unit to mark respective bits in the bitmap corresponding to the locations that are included in the record.

60. The product according to claim 41, wherein the second storage subsystem comprises a computer workstation, and wherein the product comprises an installable file system, which causes the second control unit to maintain the record and to provide the data to the host processor.

61. A method for managing a data storage system that includes first and second storage subsystems, the method comprising:

15 copying data stored on the first storage subsystem to the second storage subsystem in an asynchronous mirroring process;

maintaining a record on the second storage subsystem, indicative of locations at which the data have been updated on the first storage subsystem and have not yet been copied to

20 the second storage subsystem;

receiving at the second storage subsystem, from a host processor, a request to access the data stored at a specified location on the data storage system; and

if the specified location is included in the record,

25 initiating a synchronous transfer of the data at the specified location from the first storage subsystem.

62. The method according to claim 61, and comprising providing the data stored at the specified location from the second storage subsystem to the host processor in response to the request without initiating the synchronous transfer if the
5 specified location is not included in the record.

63. A method for managing a data storage system that includes first and second storage subsystems, the method comprising:
storing data on both the first and the second storage
10 subsystems;
receiving first updates to the data at the first storage subsystem, and receiving second updates to the data at the second storage subsystem; and
performing a symmetrical mirroring process, so as to copy
15 the first updates from the first storage subsystem to the second storage subsystem, and to copy the second updates from the second storage subsystem to the first storage subsystem.

64. The method according to claim 63, wherein performing the
20 symmetrical mirroring process comprises copying the updates between the storage subsystems in an asynchronous mirroring process.

65. The method according to claim 63, wherein receiving the
25 first and second updates comprises receiving first and second data write requests submitted by host processors to the first and second storage subsystems, respectively.

66. The method according to claim 63, and comprising:

receiving first and second data requests submitted by first and second host processors to the first and second storage subsystems to access the data stored at first and second locations, respectively, in the data storage system;

providing the data from the first storage subsystem to the first host processor after ascertaining that the data stored at the first location on the first storage subsystem are up to date; and

providing the data from the second storage subsystem to the second host processor after ascertaining that the data stored at the second location on the second storage subsystem are up to date.

67. The method according to claim 66, wherein performing the symmetrical mirroring process comprises maintaining a respective record on each of the storage subsystems, indicative of locations at which the data have been updated on the other of the storage subsystems and have not yet been copied in the symmetrical mirroring process, and wherein providing the data from the first and second storage subsystems comprises determining, responsively to the respective record, whether to request a synchronous update of the data before providing the data to one of the host processors.

68. A data storage system, comprising first and second storage subsystems, which are arranged to store data,

wherein the first storage subsystem is arranged to copy the data stored on the first storage subsystem to the second storage subsystem in an asynchronous mirroring process, and

5 wherein the second storage subsystem is arranged to maintain a record indicative of locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second storage subsystem, and

10 wherein the second storage subsystem is further arranged to receive a request from a host processor to access the data stored at a specified location on the data storage system, and if the specified location is included in the record, to initiate a synchronous transfer of the data at the specified location from the first storage subsystem.

15 69. The system according to claim 68, wherein the second storage subsystem is arranged to provide the data stored at the specified location to the host processor in response to the request without initiating the synchronous transfer if the specified location is not included in the record.

20

70. A data storage system, comprising first and second storage subsystems, which are arranged to store data,

25 wherein the first storage system is arranged to receive first updates to the data, while the second storage system is arranged to receive second updates to the data, and

wherein the first and second storage systems are arranged to perform a symmetrical mirroring process, so as to copy the first updates from the first storage subsystem to the second

storage subsystem, and to copy the second updates from the second storage subsystem to the first storage subsystem.

71. The system according to claim 70, wherein the symmetrical
5 mirroring process is an asynchronous mirroring process.

72. The system according to claim 70, wherein the first and second updates comprise first and second data write requests submitted by host processors to the first and second storage
10 subsystems, respectively.

73. The system according to claim 70, wherein the first and second storage subsystems are arranged to receive respective first and second data requests submitted by first and second
15 host processors to access the data stored at first and second locations, respectively, in the data storage system, and

wherein the first storage subsystem is arranged to provide the data to the first host processor after ascertaining that the data stored at the first location on the
20 first storage subsystem are up to date, and

wherein the second storage subsystem is arranged to provide the data to the second host processor after ascertaining that the data stored at the second location on the second storage subsystem are up to date.

25

74. The system according to claim 73, wherein each of the first and second storage subsystems is arranged to maintain a respective record indicative of locations at which the data

have been updated on the other of the storage subsystems and have not yet been copied in the symmetrical mirroring process, and to provide the data after determining, responsively to the respective record, whether to request a synchronous update of the data.

75. A computer software product for use in a data storage system including primary and secondary storage subsystems, which are arranged to store data and include respective first and second control units, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the first and second control units, cause the first control unit to copy the data stored on the first storage subsystem to the second storage subsystem in an asynchronous mirroring process, and

wherein the instructions cause the second control unit to maintain a record indicative of locations at which the data have been updated on the first storage subsystem and have not yet been copied to the second storage subsystem, and further cause the second control unit, upon receiving a request from a host processor to access the data stored at a specified location on the data storage system, if the specified location is included in the record, to initiate a synchronous transfer of the data at the specified location from the first storage subsystem.

76. The product according to claim 75, wherein the instructions cause the second control unit to provide the data stored at the specified location to the host processor in

response to the request without initiating the synchronous transfer if the specified location is not included in the record.

5 77. A computer software product for use in a data storage system including primary and secondary storage subsystems, which are arranged to store data and include respective first and second control units, the product comprising a computer-readable medium in which program instructions are stored,
10 which instructions, when read by the first and second control units, cause the first control unit to receive first updates to the data, and cause the second control unit to receive second updates to the data, and further cause the first and second control units to perform a symmetrical mirroring
15 process, so as to copy the first updates from the first storage subsystem to the second storage subsystem, and to copy the second updates from the second storage subsystem to the first storage subsystem.

20 78. The product according to claim 77, wherein the symmetrical mirroring process is an asynchronous mirroring process.

79. The product according to claim 77, wherein the first and
25 second updates comprise first and second data write requests submitted by host processors to the first and second storage subsystems, respectively.

80. The product according to claim 77, wherein the instructions cause the first and second control units to receive respective first and second data requests submitted by first and second host processors to access the data stored at first and second locations, respectively, in the data storage system, and

wherein the instructions cause the first control unit to provide the data to the first host processor after ascertaining that the data stored at the first location on the first storage subsystem are up to date, and cause the second control unit to provide the data to the second host processor after ascertaining that the data stored at the second location on the second storage subsystem are up to date.

81. The product according to claim 80, wherein the instructions cause each of the first and second control units to maintain a respective record indicative of locations at which the data have been updated on the other of the storage subsystems and have not yet been copied in the symmetrical mirroring process, and to provide the data after determining, responsively to the respective record, whether to request a synchronous update of the data.